

Inside this issue

Feature Articles

Modeling Ultra-Low Loss Accelerators	1
Advanced Networking Projects Support High-Performance Computing at Los Alamos	2
Capturing E-mail as a Record at LANL	5
New Weekly Alerting Service via SciSearch at LANL	6
The Phasing Out of Beta and Its Alternatives	8
PAGES Replaces ILFORD Printer with FUJI Printer	9

WWW at LANL

Getting a Web Site Indexed	10
----------------------------	----

Tips from the Consultants

TeX on the Cluster	13
--------------------	----

Cluster Corner

IBM XL High Performance Fortran Now Available on the Open Cluster	14
---	----

In the Classroom

Research Library Training	17
Labwide Systems Training	18
Vendor Computer Training	20

Index	27
-------	----

Using the Advanced Computing Laboratory's CM-5 computer, scientists at LANL are developing parallel simulation codes to perform high-resolution modeling of charged particle beams in accelerators. The figure above depicts simulation results that describe the transport of an intense charged particle beam in a uniform focusing channel. From top to bottom it shows snapshots of the initial, intermediate, and final states of the beam; the left column shows the horizontal phase space (x, p_x), and the right column shows the beam in real space (x, y). If the beam were properly matched to the transport system, the results would always look similar to the starting plots, as shown in the top frames. But because of an instability in this type of beam (known as a Kapchinskij-Vladimirskij or KV distribution), particles break away from the core and form a large-amplitude halo, as shown in the lower frames. This simulation was developed and run by Robert Ryne of the Accelerator Operations Technology Division and Salman Habib of the Theoretical Division. For more information, see the article on page 1 in this issue.

CIC Customer Service Center(505) 665-4444 or cichelp@lanl.gov

Integrated Computing Network (ICN)

Consulting:

Centralized scientific and engineering computingconsult@lanl.gov or 7-5746

Lab-wide administrative and business systems.....labwide@lanl.gov or 7-9444

Passwords (required for access to ICN)validate@lanl.gov or 5-1805

Central Computing Facility (CCF)7-4584

Advanced Computing Laboratory (ACL)5-4530

Desktop Support Center (DSC)7-4357 (7-HELP)

For PC questions: PC-help@lanl.gov or 7-9372

For Macintosh questions: Mac-help@lanl.gov or 5-1361

For UNIX questions: UNIX-help@lanl.gov or 5-2220

For groups with CIC-2 support contracts: 5-2220

Telephone Services Center7-3400

(includes voice mail)

Computer training

Lab-wide systems support training7-9444

Computer/workstation training7-9399

Personal computer training7-9071

Microcomputer support facility seminars7-4357

(Macintosh/IBM software, lending library)

Network Operations Center (NOC).....noc@lanl.gov or 7-7423

(after hours call 7-4585)

List of Forms

ICN Validation Request Form23

Reader Feedback Form25

Modeling Ultra-Low Loss Accelerators

The United States is now involved in efforts aimed at developing accelerator-driven technologies for transmutation of radioactive waste, disposal of plutonium, energy production, and tritium production. These technologies have highly beneficial implications for the nation's environment, security, and future energy needs. Similar accelerator technology will also have a major impact on the U.S. scientific infrastructure through the development of next-generation spallation neutron sources for materials science and biological science research.

All of these projects require accelerators that operate with extremely low beam loss: at the high-energy end the allowed loss is less than one billionth of an ampere per meter. It is now known that a major source of beam loss is the formation of a very low density halo far from the beam core. Understanding and predicting beam halo and finding ways to minimize it will have a significant impact on the above-mentioned technologies. In order to reduce cost and technological risk, and to improve accelerator efficiency, performance, and reliability, it will be necessary to perform high-resolution modeling far beyond that which has ever been performed by the accelerator community.

Using the resources of the Advanced Computing Laboratory at LANL, we are developing tools that will enable us to model the dynamics of intense charged particle beams with unprecedented speed and accuracy. Previously, most linear accelerators were designed and optimized using particle simulation codes that were run with 10,000 to 100,000 particles. But in the new ultra-low loss regime, predicting the beam halo with confidence will require simulations that use on the order of 100 million particles. Though such simulations would have been impossible in the past, advances in computer hardware and algorithms are making such high-resolution simulations feasible.

Through a joint effort of the LANL Accelerator Operations and Technology Division, the Theoretical Division, and the Advanced Computing Laboratory, we are developing accelerator simulation codes for high performance computing platforms. The codes provide us with a means to test and optimize new accelerator designs and test theoretical models. For example, we have performed multi-million particle simulations that have confirmed the validity of a new theoretical model called the particle-core model, which describes halo evolution in mismatched charged particle beams. Along with developing particle simulation codes, we are also studying other approaches to modeling and predicting beam halo. For example, the large memory available

on massively parallel computers makes it possible to utilize direct Vlasov-Poisson solvers. We have developed a two-dimensional Vlasov-Poisson code that was run on the CM-5 with 128 grid points in each of the four phase space dimensions for a total of 268 million grid points.

Though the main emphasis of this work is accelerator simulation, the tools we are developing are applicable to other areas. For example, our particle simulation codes have been modified to study the formation of large-scale structure in the universe. Also, our direct solver techniques have been applied not only to the Vlasov-Poisson system but also to quantum systems described by Schrodinger wave functions, density matrices, and Wigner distribution functions. Using these codes we have performed simulations of mesoscopic systems on the CM-5 in a matter of hours that would have taken months to run on a workstation.



This work is supported by the U.S. Department of Energy, Office of Energy Research, through the Division of Mathematical, Information, and Computational Sciences, and the Division of High Energy Physics.

*Robert Ryne, ryne@lanl.gov, (505) 667-2839
Accelerator Operations and Technology Division*

*Salman Habib, habib@lanl.gov, (505) 667-5265
Theoretical Division (T-8)*

Advanced Networking Projects Support High-Performance Computing at Los Alamos

High-performance networks can be defined as computer networks that support high-performance computing. Before 1988, a high-performance network was thought to be 10 megabit/second Ethernet with the potential for 100 megabit/second FDDI (fiber distributed data interface) on the horizon. In 1988, Los Alamos started development of an 800 megabit/second network link. The catalyst for this effort was the requirement for visualization of computer simulations at resolutions and data rates much higher than the computer networks of the late 1980s could support. High-resolution visualizations (at least 1000 x 1000 picture elements, or pixels) provide scientists with more detail and allow them to view consecutive frames of a simulation in rapid succession.

Now, in 1996, the requirement for high-performance networks is being driven from a new source: supercomputers based on symmetric multiprocessor (SMP) computing systems. These new systems have a much larger networking component than traditional supercomputers. Also, with the new SMP systems come new applications that depend on the SMP and its interconnecting network to provide the high performance environment needed to run these applications. Today, networks running at nearly one gigabit/second are considered legacy high-performance networks that can't meet these new requirements, so the push is on to develop multi-gigabit networks that will.

Setting The Stage For Gigabit Network Research: The High Performance Parallel Interface (HIPPI)

In 1987, network researchers in the Network Engineering Group (CIC-5) began developing HIPPI, which is a computer interface that operates at 800 megabits/second, the rate needed to support high-resolution visualization. In addition, it was recognized that such a link could also be used to shorten the time to transfer data between computers and between a computer and

an archival storage system. As larger and faster supercomputers were being deployed at Los Alamos, application programs required larger amounts of data. Today, some applications require access to gigabytes of data that are stored elsewhere on the network. Lowering the time to transfer this data by a factor of 10 would significantly improve the time the application needs to complete its task.

One of the main goals for HIPPI was to make it a national standard so that Los Alamos engineers would not have to design a special interface for each new supercomputer the Laboratory procured. This effort was successful in 1991 when HIPPI became an American National Standard Institute (ANSI) standard. HIPPI is now the gigabit/second standard of choice for most supercomputers and high-end workstations. It is available on computers from Cray, IBM, Thinking Machines, Intel, SGI, Hewlett-Packard, Sun, and many other computer vendors.

Los Alamos researchers then developed a network based on HIPPI that used crosspoint (crossbar) switches to increase the aggregate bandwidth of the overall network. It is possible to interconnect several HIPPI crossbar switches into a multiple crossbar network. In addition, gateways to existing legacy networks, such as Ethernet and FDDI, are now available. Figure 1 shows such a crossbar network.

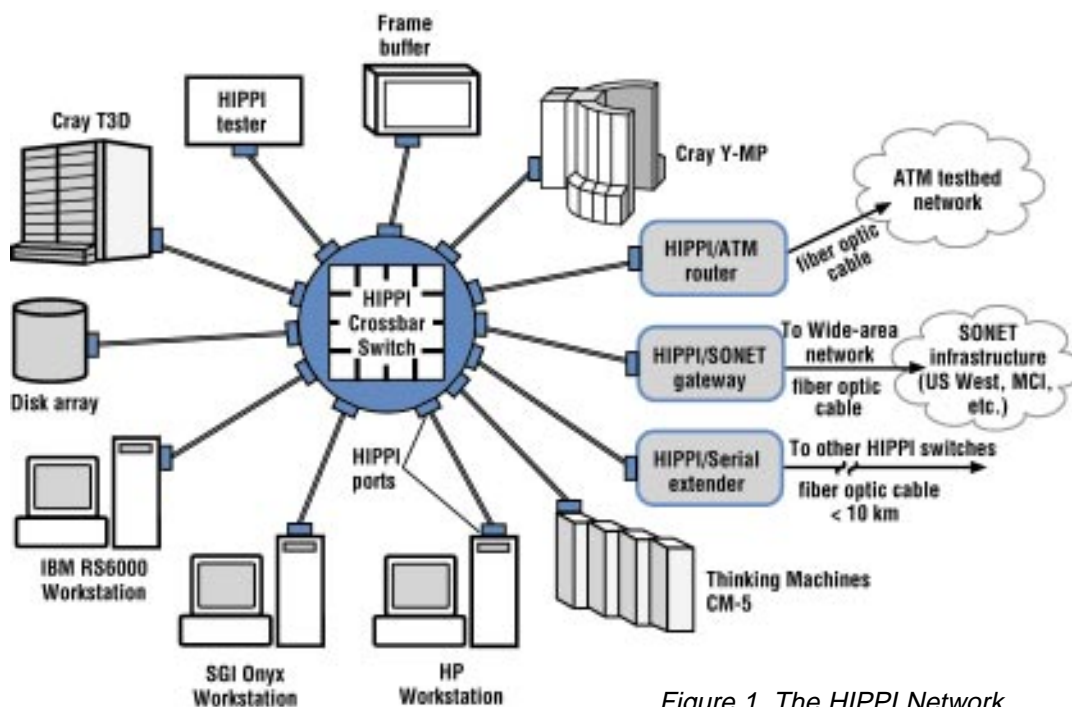


Figure 1. The HIPPI Network

The HIPPI network at Los Alamos is much larger than the one shown in Figure 1. The current network now has six crossbar switches interconnecting supercomputers in the Advanced Computing Laboratory (ACL), the Central Computing Facility, and the Laboratory Data Communications Center. Much of the HIPPI network is now in production status. CIC-5 has worked with personnel from the Computing Group (CIC-7) and the Storage Group (CIC-11) to migrate Central File System (CFS) data links to the supercomputers from the existing FDDI backbone to a HIPPI backbone. This migration occurred in the secure partition first (February 1996) because the demand for large data transfers between CFS and the supercomputers is higher than in the open partition.

CASA Gigabit Testbed

The Laboratory's success with HIPPI in 1988-90 brought an offer to participate in the CASA Gigabit Testbed (one of five such testbeds). This effort was started in 1991 and was jointly funded by the Advanced Research Project Agency and the National Science Foundation. The CASA Testbed developers took the Los Alamos HIPPI network design and extended it to each of the four CASA sites: Los Alamos, San Diego Supercomputing Center (SDSC), Caltech, and Jet Propulsion Laboratory (JPL).

The main challenge for Los Alamos was to develop a technique to interconnect the HIPPI crossbar networks at each site at the full 800 megabit/second data rate. Toward this end, the telephone carriers have been installing, between the sites, high-capacity lines based on another standard called synchronous optical network (SONET). Los Alamos engineers developed the HIPPI/SONET Gateway to provide an interface between HIPPI and SONET. This unit was deployed into the testbed in 1994, and the link between Los Alamos and SDSC became operational in July 1994. The HIPPI/SONET Gateway won a national 1995 R&D 100 Award for being one of the top 100 technological innovations of the year. Figure 2 shows the location of the CASA Testbed and the telecommunication carriers who donated the SONET links.

The CASA Testbed effort ended in 1995. While it was operational, several benchmark tests were done on the long distance portions of the network. CASA has the longest links of all five gigabit testbeds; thus, it has the largest amount of data being carried on the links and, as a result, the greatest latency. This gave CASA some special characteristics that are useful for network research. In December 1994, testing was done over the link with the TCP/IP protocol suite that is used in the Internet.

Although protocols over networks ensure data-transfer integrity, they also lower the effective data rate. However, even with this effect, the data transfer rate achieved for TCP/IP over CASA between San Diego and Los Alamos was a record setting 550 megabits/second.

Other Gigabit Efforts

Figure 1 also shows a connection to an asynchronous transfer mode (ATM) testbed. ATM is a new networking technology being promoted by the telecommunications industry. ATM has the attractive feature of supporting voice, video, and data rates at very high speeds, with a basic data-transfer rate of 155 megabits/second. What makes ATM unique is the use of small 53-byte cells to transfer all information. ATM also scales up into gigabit/second data rates when coupled with SONET as its physical transfer mechanism. Although ATM offers several attractive features, it is not clear whether this technology is suitable for supercomputer applications.

Los Alamos has set up an ATM testbed to evaluate ATM networks and had a cooperative research and development agreement (CRADA) with Digital Equipment Corporation in 1995 to develop an ATM frame buffer for visualization over ATM networks. Los Alamos engineers developed this frame buffer using eight clustered workstations, each sending one-eighth of a high-resolution image (1,000 x 2,000 pixels). The aggregate rate for transferring data from all eight workstations to the frame buffer is about 1.4 gigabits/second. The successful operation of this system paves the way for using clustered workstations for other such data transfer applications at Los Alamos and other DOE laboratories.

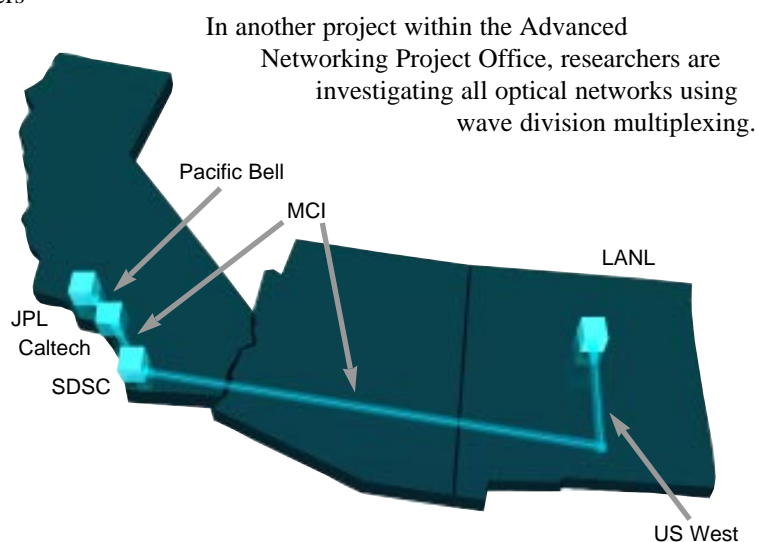


Figure 2. The CASA Testbed

This effort was started as part of a CRADA with IBM to upgrade an IBM-developed optical network to run at gigabit/second rates. HIPPI was used as the link between computers and this network because it was the only gigabit/second standard available at the time (1992). What makes such networks attractive is that they allow more of the inherent bandwidth to be used in each fiber-optic cable, and they eliminate the need to use an expensive and potentially unreliable (compared to pure optics) electrical switch in the middle. The passive optical switch used in the network is an inexpensive optical coupler that sends all wavelengths to each destination. The filter at each end node then picks out the wavelength intended for that node.

Now that this CRADA is finished, additional work is in progress to lower the latency of optical filter switching to achieve an efficient packet-switched optical network. Current filters switch in several milliseconds. The goal is to get this down to several microseconds to more efficiently support gigabit transfer rates. Interest in this technology is growing as the requirements on bandwidth go beyond the 10 gigabit/second barrier. Optical networks may be the only way to achieve data rates above this level.

Development of Next-Generation Networks

Gigabit networks have now become a mature technology, with HIPPI being deployed at most supercomputer centers around the world. The Advanced Networking Project in CIC Division is now establishing direction for the next generation of high-performance networks. The requirement for data rates approaching 10 gigabits/second is growing out of the Accelerated Strategic Computing Initiative (ASCI) and application requirements for the Laboratory's High Performance Computing Research Center (HPCRC), which is part of the ACL. Because both programs have identified SMP computing systems as their future core compute resources, the network demands for such systems has given rise to our effort to develop the next generation of HIPPI networks.

These networks have both a bandwidth requirement and a latency requirement. The bandwidth must be multigigabits/second and the latency must be in the few microsecond region. The latency issue is the most difficult to deal with because it depends on other constraints besides the physical network. Such things as protocol processing, operating system interaction, and retransmission of erroneous data severely impact latency. Therefore, the approach being taken is to develop a network that deals with some of these issues directly at the physical layer of the network. The result is evolving to a new version of HIPPI called HIPPI-6400. Features of HIPPI-6400 include the following:

- 6.4 gigabits/second bandwidth,
- Automatic retransmission of erroneous packets,
- Fixed size micro-packet (32 bytes) transfer unit,
- Virtual channels to avoid congestion (and improve latency), and
- Credit based flow control to prevent buffer overflow.

Los Alamos network developers are currently leading the effort to standardize HIPPI-6400 through the American National Standards Institute (ANSI). They are also building network components and testers in collaboration with companies which will provide the necessary network infrastructure to support this new high-performance multigigabit network.

*Stephen Tenbrink, sct@lanl.gov, (505) 667-4935
Network Engineering Group (CIC-5)*

Capturing E-mail as a Record at LANL

If you use e-mail for your work at the Lab, you have probably struggled with a common problem: How much of the volume of messages is really worth saving? In your specific area of work, you may have a good sense of how much, as well as for how long they should be kept. However, in many areas of LANL activity, the significance of e-mail as a crucial part of the record that needs to be preserved is many times overlooked by its authors.

In addition, e-mail is coming under stricter regulations. Effective October 1, 1997, all federal agencies and their contractors are required by law to have a system in place to capture E-mail Record Material. Here are the main driving forces behind this new concern:

- Increased usage of e-mail within LANL and the DOE
Complex to communicate and document decisions, and the accompanying fear that significant Department records are not being preserved because they exist only in electronic form.
- Recent legal decisions, new law, and DOE guidance.

Increased Usage. The key concern behind increased usage and fear of loss is just good business practice. Increasingly, the record of decisions that affect the operation, safety, or business aspects of LANL take place only on electronic media, and the proper filing and retention of e-mail, much less hard copies, is not always a given.

Legal. The conclusion of a major federal court case was that government agencies must retain certain electronic documents as records. In the 1993 *Armstrong vs. Executive Office of President* (a follow-on to the noted Ollie North electronic mail case), the court also ruled that agencies must retain Header and Recipient information (e.g., what did they know and when did they know it), and that users must be trained in the requirements for handling electronic records. The resultant DOE mandates also apply to other electronic processes (such as STU-III secure teletypes), but e-mail carries the most information at this time and is the most prevalent medium involved.

The Information and Records Management Group (CIC-10), along with other CIC partners, is evaluating different technical implementations that identify e-mail when it is of "record" quality. In the long-term, when commercial software becomes available, an inference engine will automatically identify an e-mail message as a record (if applicable), its record type, its retention, and then index the record copy. This type of technology would be integrated with the current efforts in developing an institutional records management process and system.

What E-mail is a Record?

Certainly not all e-mail is of record quality, according to Tom Reding, Deputy Group Leader for CIC-10. The significance usually lies in those messages that provide documentary evidence of the business operations of the organization or that have permanent historical value. Reding's organization includes the Lab Archivist and Historian who is available to help establish criteria for historical value.

Examples of record-quality types of e-mail include:

- All e-mail to or from members of the Laboratory Leadership Council;
- E-mail generating or exchanging LANL financial, contractual, or resource commitments;
- E-mail that contains legal information or that deals with requests through Freedom of Information Act, California Public Records Act, California Information Practices Act, or other litigation;
- E-mail that relates to Environmental, Health, or Safety issues or that responds to DOE directives;
- E-mail to or from Congressmen, Senators, or DOE Managers; and
- E-mail about security infractions, audits, appraisals, surveys, inventories, foreign trip reports, action plans, or proposals.

These e-mail guidelines are identical to the ones CIC-10 has already developed for traditional "hard copy" paper correspondence managed by the Correspondence Management Team, also operated by CIC-10. Even if there is a hard copy of an electronic memo or attachments, e-mail transmission and recipient information is a key legal issue. Reding's group also specializes in classified correspondence control, and special consideration will be given to that category of e-mail, he added.

Meeting the October 1, 1997, deadline is critical for LANL to demonstrate its commitment to the Administration, Secretary O'Leary's Openness Initiative, and integrity in records management, Reding emphasized.

For further information, call the CIC-10 Group Office at 7-5330 or send e-mail to treiding@lanl.gov.

*Bill England, bengland@lanl.gov, (505) 665-0208
Communications Arts and Services (CIC-1)*

New Weekly Alerting Service via SciSearch at LANL

The Los Alamos Research Library (CIC-14) would like to introduce a new literature searching program called Weekly Alerting Service, which is one of the options available via SciSearch at LANL. Weekly Alerting Service allows researchers to design their own search strategy and then automatically receive notification of newly published material relative to their own area of expertise or interest. Access to Weekly Alerting Service is available via SciSearch at LANL, an electronic index covering thousands of science and technology journals. To enter SciSearch at LANL, point your Web browser at

<http://bighorn.lanl.gov:4001>

Before using the Weekly Alerting Service, we recommend that you test your search criteria in the SciSearch at LANL database as either a General Search or a Cited Search. This will allow you to look at the search results and then determine whether you need to broaden or narrow your search criteria.

To use Weekly Alerting Service, you must first register and login. From the Web page cited above, select "Alerts" to access the main Weekly Alerting Service Web page. Now select "Register" and then enter your name, Z-number, a password of your choosing, and your e-mail address. After registering, select the "login" option and enter your name and password. Now you're ready to input the criteria for your automatic literature search.

Weekly Alerting Service provides two types of literature searches: General and Cited. The general search provides a listing of recent article titles based on the criteria you provide, i.e., author, keyword, or journal (see Figure 1). The cited search allows you to enter an author's name, Albert Einstein for instance, and then receive a listing of recent articles that have cited articles

written by your chosen author, in this case articles by Albert Einstein (see Figure 2). Once you have entered the criteria for your search, Weekly Alerting Service will automatically conduct your literature search and notify you of the results on a weekly basis.

Notification of search results is accomplished by one of three possible options: Notification only, Citations, or Citations and Abstracts (see Figures 1 and 2). If you select Notification only, you will receive an e-mail indicating the number of "hits" turned up by the search. Then you can login to

Search **Home** **Next** **Comments** **Mark** **Library Databases**
Up **Previous** **Help!** **Download**

Create a General SciSearch® at LANL Alert

Specify options:

Alert Name

Email Address

Receive ☒ **Notification Only** ☐ **Citations** ☐ **Citations and Abstracts**

Specify search criteria:

Author

Title Words

Abstract, Title or Keywords

Source (Exact match ☐)

Institution

Categories

All Fields

Maximum number of citations to retrieve: **Submit** **Clear**

Figure 1. General SciSearch Window

SciSearch (<http://bighorn.lanl.gov:4001>) and see exactly what articles were found. This option allows you to take advantage of hyperlinks and other features provided by SciSearch. If you select the Citations option of receiving notification, you will receive an e-mail containing the actual titles of articles turned up by the search. Selecting Citations and Abstracts provides article titles and, when available, abstracts.

CIC-14 also offers a free class called SciSearch Alerting Service. This class provides basic instruction and an overview of the many additional features that were not covered in this

article. The class will be offered November 7 at 1:30 p.m. in the Research Library. For more information about the class contact the Research Desk at 7-5809 or library@lanl.gov. For general questions about Weekly Alerting Service, call the Research Library at 667-5809 or send e-mail to library@lanl.gov.

*Frances Knudson, fknudson@lanl.gov, (505) 667-3031
Librarian / Research Library (CIC-14)*

The screenshot shows a web interface for creating a SciSearch alert. At the top is a navigation bar with buttons: Search, Home, Next, Comments, Mark, Library Databases, Up, Previous, Help!, and Download. Below this is the title "Create a Cited SciSearch® at LANL Alert". The form is divided into two main sections. The first section, "Specify options:", contains fields for "Alert Name" and "Email Address" (with "finney@lanl.gov" entered). Below these are radio buttons for "Receive": "Notification Only" (selected), "Citations", and "Citations and Abstracts". The second section, "Specify criteria for cited reference:", includes a note "(First Author's last name and initials are required.)" and fields for "First Author's Last Name", "First Author's Initials", "Year", "Source (1st char)", "Volume Number", and "First Page Number". At the bottom, there is a field for "Maximum number of citations to retrieve:" set to "50", and "Submit" and "Clear" buttons.

Search	Home	Next	Comments	Mark	Library Databases
Up	Previous	Help!	Download		

Create a Cited SciSearch® at LANL Alert

Specify options:

Alert Name:

Email Address:

Receive: ☒ Notification Only ☐ Citations ☐ Citations and Abstracts

Specify criteria for cited reference:
(First Author's last name and initials are required.)

First Author's Last Name:

First Author's Initials:

Year:

Source (1st char):

Volume Number:

First Page Number:

Maximum number of citations to retrieve: ▼

Figure 2. Cited SciSearch Window

The Phasing Out of Beta and Its Alternatives

Effective December 1, 1996, services on machine Beta will no longer be available. Because Beta was used for a variety of functions, the Beta team did not want to just “pull the plug” without providing suitable alternatives and offering assistance during this transition period.

In August 1996, the ICN administrators surveyed all Beta users and asked them about their dependency on Beta. There were 104 individual responses (about 10% of the total number of Beta users). The majority indicated that they never or seldom log-on to the machine. Those who logged on daily overwhelmingly indicated that they did so only for e-mail. The survey also indicated that e-mail was the number one overall use for the machine. To respond to the need for an alternative e-mail service and other needs revealed by the survey, the Beta support team came up with the following eleven Beta functions and some viable alternatives.

1. E-mail. If you use Beta for UNIX mail, you will need to set up UNIX mail on your individual network or UNIX workstation. PINE and mailx (use mailx if you used to use mail) is also available on all of the CIC e-mail (POP) servers. Your local system administrator can help you with that. If you use POP e-mail(i.e., Eudora) on Beta, CIC maintains three POP e-mail servers. These machines allow the use of POP clients (for example, Eudora) from anywhere, including TCP/IP connections and dialup. CIC's POP clients are cheaper for the paying customer. A Beta account costs \$45.00 a month while a POP account only runs \$15 a month. POP saves each customer \$360 a year! You can register for a POP account by telnetting to register.lanl.gov or by contacting the Customer Service Center at (505) 665-4444 option 1 or cichelp@lanl.gov.

2. Internet access. Users were first telnetting to Beta and then telnetting to the rest of the world. There is no need to do this. If you can get to Beta, you can get to the entire Internet. This is safer now with the use of a Smart Card for remote log-on. To get out of the administrative partition and into the open Internet, you will first need to telnet to adgate.lanl.gov. From there you can telnet through the fire wall.

3. File Storage. Some users indicated that they were storing files either short term or long term on Beta. The Cluster or UNICOS machines will store large files for 24 hours in /usr/tmp. The CFS is available for long term storage.

4. Usenet News. There is a plethora of usenet clients for all machines. Netscape 2+, a LANL IA standard, also functions as a news reader. Netscape is available for Cluster Users.

Terminal Emulating users can run PINE on the CIC POP servers to read news. Alternatively, trn is on the Open Cluster.

5. Programming. Some users wrote and maintained C and Fortran codes on Beta. The exact same compilers are available on the Open Cluster.

6. File Transfer. Some users indicated that they used Beta to ftp files around the Internet. FTP is available for all platforms and is standard on the Cluster and UNICOS. To get out of the administrative partition and into the Open Internet, you will first need to ftp to adgate.lanl.gov. From there you can ftp through the fire wall. Some GUI ftp clients (like the one provided with Windows 3.1) will not work properly through adgate. The Macintosh product, Fetch, will work through adgate if you issue ftp commands from the pull-down menus. In addition, kermit is available on the Open Cluster and the mail servers.

7. TeX/LaTeX. These utilities have been installed on the Cluster. See the BITS article titled “TeX on the Cluster” in this issue for more information.

8. PAGES. Some users indicated that they submitted jobs to PAGES from Beta. PAGES access is available from many platforms and can be mounted on individual machines. The Cluster and UNICOS machines have PAGES access as a standard.

9. Lynx. All lynx users can telnet to www.lanl.gov 1664. Note that the port is 1664. After log-on, select option 3. lynx v 2.3A.

10. FRED. The FRED editor is standard on UNICOS and the Cluster.

11. tar. The ability to tar files for long time storage is available on all UNIX machines including UNICOS and the Cluster.

If you still need assistance with the transition from Beta, please contact the ICN Consulting Office at 667-5746 or consult@lanl.gov.

*Dale Hugo Leschnitzer, consult@lanl.gov, (505) 667-5746
Customer Service Group (CIC-6) / ICN Consulting Office*

PAGES Replaces ILFORD Printer with FUJI Printer

The PAGES Team will remove the ILFORD digital photo imaging printer from service on Friday, November 22, 1996, at 10:00 a.m. This device is being removed from service because the vendor has gone out of business and will no longer provide a source of replacement parts for the device.

In anticipation of its removal, the PAGES Team has already introduced the FUJI Pictography 3000 printer on a "friendly user" basis. These friendly users, many of whom are our traditional ILFORD customers, are more than satisfied with the quality of the output coming off of the FUJI. To run a job on the FUJI, customers will not have to do anything different than they are already doing. Simply continue to use the sheet photo or trans photo options in the PPAGES or PSPAGES shipping utilities.

The FUJI service is already available in the Open partition, and it will be available in the Secure partition by November 22.

Users who submit sheet photo or trans photo jobs in the Secure partition before November 22 will receive output from the ILFORD device. On November 22, the operations staff will simply begin running the jobs (sheet photo or trans photo) to the FUJI device.

The FUJI device is very similar to a dye-sublimation type device. It offers a near-photographic type of imaging. The Fuji is a step up from the current sheet/trans color (QMS wax-transfer type) system in that it provides a higher quality image (400 dpi versus 300 dpi) and a larger image area (7.9" X 10.5" versus 7.9" X 8.75").

If you have any questions regarding this change in services, please contact the PAGES operator at 667-2905 or me personally at the number listed below.

*Reuben Roybal, rrr@lanl.gov
(505) 665-2074
Media Group (CIC-17)
PAGES Team Leader*



Getting a Web Site Indexed

Imagine a library. A big one. Ten or fifteen city blocks in area, and about as many floors. Go through the library, pull all the books from the shelves, and mix them all together on the floor. To that mix, add a bunch of diary entries, postcards, photo albums, loose ramblings, and whatever else you can think of.

Then stand back, look at it—that massive collection of information about anything you might want and plenty that you don't—and try, just try, to find the piece of information you want.

Sometimes (too often) that still seems to be the state of the Web. I've heard it compared to libraries before the Dewey Decimal System: plenty of information, but no clear path for finding what you're looking for. (This is, I think, the root source of the complaint that browsing is a non-productive activity—which is, of course, an assessment I disagree with.)

On the positive side, the quality of indexing mechanisms continues to improve as their number grows. In addition to some excellent work done here at the Laboratory (see, for example, the INDEX project), industry solutions have advanced from site-level cataloguing (e.g., Yahoo) to increasingly sophisticated page-level indexes (e.g., AltaVista or HotBot, each of which boasts more than 40 million pages indexed).

Although these and rival tools still have plenty of room for improvement (frequency of updates comes immediately to mind), they already offer Web administrators some valuable tools for helping others find what they're looking for within our Web space, for drawing more attention to our Web space, and for maintaining the Web space itself.

This article addresses some of the practical ways we can make use of the indexing tools to improve our Web space. It won't address query syntax or evaluate "which tool is better," but it will focus on ways to leverage the work done by others to let the tools work for us.

Ready or Not, You're Getting Indexed

Modern indexing tools use software "robots" that automatically go out, search the Web, and report their results to a central index. (If you monitor your access logs, most of these robots are easily recognized as machine names that run up a large number of hits with only a few seconds between each hit—not enough time for a real person to actually read the pages.) Unless you have specifically excluded these robots, any publicly accessible Web space you've constructed has probably already been indexed. If the Web space wasn't yet ready, the results are likely to be somewhat less flattering than you might hope.

For example, an AltaVista search of the Information Architecture project's Web space turns up, among other things, the following entry for one of our standards:

IA-6101: GIS Data Transfer—SDTS
IA Home | Laboratory Home | Introduction | What's New | IA Glossary
Standards | RFCs* | Teams* | Areas* | IA Library | Navigational Aids.
Strategic...
<http://www.lanl.gov/projects/ia/stds/ia610110.html> - size

Figure 1. Outdated IA Web Space as Seen by AltaVista

It's easy for me to recognize that all the garble in the description is the old navigation bar we were using back in March 1996—but it's not very helpful to anybody else who might be looking for information about the transfer of spatial data (which is what the standard addresses).

For comparison, here's the result for a more recently adopted standard:

IA-6402: Guidelines for HyperText Links from Laboratory WWW
Guidelines for links to commercial, non-profit, government, and Laboratory-community sites.
<http://www.lanl.gov/projects/ia/stds/ia640212.html> - size 10K - 23 Sep 96

Figure 2. Updated IA Web Space as Seen by AltaVista

In this case, I was ready for the robot when it came and provided it with the information it needed to make a meaningful description of the material.

Another point illustrated by Figure 1 is that the indexes do not, yet, get updated as frequently as we might like. The SDTS standard document has been updated for months now so that it will yield better results if the robot ever comes back. As of this writing, though, it's still the seven month-old entry that's in the index.

Note: Many robots, including AltaVista and HotBot, advertise a "Submit URL" feature that is supposed to get a document more rapidly indexed. While this can be useful for new pages, I have not found it useful for getting old entries updated. It seems we just need to wait for the robots to evolve further before updates will happen more frequently.

Document Title, Description, Keywords

Most robots look in an HTML document's <HEAD> section to determine what the document's title is, and then read through the <BODY> section to look for keywords and a reasonable-seeming description. What seems reasonable to a robot, of course, is not necessarily meaningful to anyone else.

To help improve these results, a growing trend among indexing tools is to have the robot look into the <HEAD> section for additional instructions the author may have provided.

These instructions are marked in <META> tags, part of the Internet Engineering Task Force (IETF) HTML 2.0 specification, which enable authors to specify a "NAME" and "CONTENT" for a variable. The variable is placed inside the tag, so there is no ending tag. (The tags can also be used to specify HTTP equivalents, but that use isn't applicable here.)

The three <META> variables that are gaining the widest acceptance for indexing robots are "DESCRIPTION", which provides a brief description of the document; "KEYWORDS", which lists words that should be given a higher weight in searches; and "ROBOTS", which can contain access permissions. Not all robots read all the variables, but their numbers are growing.

For the example in Figure 2, the document <HEAD> is constructed as follows (like all HTML markup, the capitalized words in the example can be in either upper or lower case):

```
<HEAD>
<BASE HREF="http://www.lanl.gov/projects/ia/stds/
ia640212.html">
<TITLE>IA-6402: Guidelines for HyperText Links from
Laboratory WWW
</TITLE>
<META NAME="DESCRIPTION" CONTENT=
"Guidelines for links to commercial, non-profit, gover
ment, and Laboratory-community sites.">
<META NAME="KEYWORDS" CONTENT="WWW,
World Wide Web, HTML, link, linking, WWW authors and
publishers">
<META NAME="ROBOTS" CONTENT="ALL">
</HEAD>
```

When a robot such as HotBot reads the above <HEAD> section, it

- Ignores the <BASE> tag (which is described in the October 1996 BITS article, "Mariachis Weave Beethoven? More Tips and Tricks");

- Reads <TITLE> for the title of the document (the first line in the search result);
- Reads the <META> "DESCRIPTION" variable for the document's description (the second line in the search result);
- Reads the <META> "KEYWORDS" variable for words that should be given a higher weight in searches (which doesn't appear in the search result); and
- Reads the <META> "ROBOTS" variable to determine whether to index the document ("ALL" means yes, which is what the robot does if the variable doesn't appear; "NOINDEX" means no, which will stop compliant robots from adding the page to their index).

The result is shown in Figure 2, which shows the title and description as I specified them (as opposed to the example in Figure 1, where the first words in the document <BODY> were used as the garbled description).

Note: Keep descriptions brief. Indexes will generally only use the first 150 or so characters; anything longer gets truncated.

Controlling What Gets Indexed

Although getting our information indexed makes it more accessible to others, there are a number of reasons why we might want certain areas to not be indexed. Some areas might contain out of date archives, others might contain temporary working drafts, some might be intended for a limited audience only, etc.

The <META> "ROBOTS" variable described above is the most convenient and flexible, but also the weakest, way to control what the robots index. It offers file-level control over access, but as of this writing, only a handful of robots read it; most still ignore it.

A stronger mechanism is a robots.txt file, as described in "A Standard for Robot Exclusion" (<http://info.webcrawler.com/mak/projects/robots/robots.html>). This file is placed in your Web site's root directory (the Web root, not the operating system root) and lists the directories you don't want robots reading:

```
User-agent: *
Disallow:      /path-1/
Disallow:      /path-2/
```

In the first line of the example above, the asterisk means "all robots." You can also list specific robots here if there are only certain ones you want to block.

In the “Disallow” lines, compliant robots will not enter any of the directories listed, or any of their subdirectories. If you’d like to reopen the directories to a specific robot without letting any others in, then start by disallowing “*” and then add a section like the following:

```
User-agent: approved-robot-name
Disallow:
```

When nothing is disallowed, then everything is open. This approach can also be used to close different directories to different robots.

The robots.txt file is obeyed by more robots than the <META> “ROBOTS” variable, but again, adherence is voluntary and not all robots understand it. (A further potential drawback is that you need permission to the Web root directory to be able to edit the file. If you’re not the system administrator, you’ll need to coordinate any updates with him/her.)

A stronger mechanism, which robots have to obey, is to use the server access configuration files (“access.conf” or “.htaccess” by default under NCSA HTTPd). For example, in the <Limit> section of the access control file, the following lines would allow everyone *except* “robot.name” to retrieve files from the corresponding directory:

```
<Limit GET PUT POST>
allow from all
deny from robot.name
</Limit>
```

The “deny” line could be repeated for any robots you wish, and those robots would be unable to index anything in that area.

The strength of this approach is that it definitely stops the specified robots. There are, however, several drawbacks: it only stops robots you’ve identified, it doesn’t differentiate between files in the same directory, sometimes robots come from machines that are shared by others who you do want to allow in (proxies, etc.), and it makes the access control files more complicated to maintain.

Other, stronger ways to limit access include password protecting a directory, removing read permission from files at the operating system level, and deleting the files altogether. All of these mechanisms work, but they all affect other users in addition to the robots.

Using Indexes to Improve Web Maintenance

All of the above discussion can be thought of as ways to help the indexing robots do their job better. In addition, however, there are ways we can use the robots to help us do our jobs.

One technique I’ve found useful is to check the indexes for everything they’ve found in one of my Web spaces. (In AltaVista, for example, search for “URL:your-url/”, or, if you want everything on your machine, “host:your.machine.name”). This can reveal out of date files that need to be updated or removed; permissions breaches that need to be repaired (i.e., files that others aren’t supposed to be able to see); and other general housekeeping that needs to be done.

Another technique is to search for sites that have linked to our Web space. (Again in AltaVista, search for “link:your-url”). Seeing how others see us is one way of measuring our effectiveness, plus this gives us the chance to notify others if their links are out of date. (Helping others keep their links up to date is more than just a courtesy; our own Web space is more effective if people following the links of others arrive at an up-to-date place.)

There are certainly other ways to leverage the work already done by the indexing robots (building a front-end tailored to our own Web site comes to mind, such as the AltaVista and other search engines that can be found under the “Search” option on the Laboratory home page, <http://www.lanl.gov/searches/searches.html>). The examples given above are simply those that I have personally found useful.

For More Information

More information about the search engines mentioned above is available at the following URLs:

- INDEX/Explorer: <http://iosun.lanl.gov:2001/explorer.html> (restricted to Laboratory machines)
- Yahoo: <http://www.yahoo.com/>
- AltaVista: <http://altavista.digital.com/>
- HotBot: <http://www.hotbot.com/>

Again, the list is only intended to be representative. More complete lists can be found at

- <http://www.lanl.gov/Internal/finding-info.html>
- <http://home.mcom.com/home/internet-search.html>
- http://www.yahoo.com/Computers_and_Internet/Internet/World_Wide_Web/Searching_the_Web/Search_Engines/

For more information about the Information Architecture Project, please visit our home page at <http://www.lanl.gov/projects/ia/> (or look under “What’s New” from the Laboratory home page). For a summary of the project’s WWW-related guidance, please select the “Web” button from our masthead. If you would like printed or e-mail copies of any of the IA materials, please contact Tad Lane at the address given below.

*Tad Lane, tad@lanl.gov, (505) 667-0886
Information Architecture Standards Editor*

TeX on the Cluster

With the demise of machine Beta close at hand, we have installed the TeX typesetting utility on the IBM Cluster of RS6000 machines. The Cluster has 10 nodes, ibm-01 through ibm-10. These nodes share the same disk space, so your home directory will be available on any node you choose to log into.

To get an account on the Cluster, follow the instructions below:

1. Telnet to the register.lanl.gov machine,
2. Sign in with your ICN user number,
3. Select option 1 to get CIC service registration,
4. Supply your ICN password or passcode,
5. Select option 2 for the ICN Compute Servers, and
6. Select 2 for AIX (ibm-cluster).

Contact the ICN Consulting Office at consult@lanl.gov or 667-5746 if you have any trouble registering for the Cluster.

TeX Version 3.14159 (C version 6.1) and LaTeX2e <1996/06/01> are the current versions on the Cluster. Fonts are generated as needed, so you may see extra messages from MakeTeXPK while it creates fonts for you.

The xdvi previewer is available if your workstation allows X window connections. xdvi allows you to preview your TeX output (with included PostScript files) on your workstation monitor.

The LANL letter and memo macros do not port to LaTeX2e. Bob Macfarlane in T-2 wrote similar LaTeX2e class macros and has allowed us to install them on the Cluster. If you use the letter or memo style files, be aware that they are incompatible with the former versions due to the

commands `\to`, `\from`, and `\symbol` already being used. These commands have been replaced with `\towhom`, `\fromwhom`, and `\referto`, respectively. Template files are available on the Cluster in the `/usr/local/texmf/tex/latex/local` directory and are called `memotest.tex` and `lettest.tex`.

The dvips utility is different. It will send the output to PAGES by default. If that is okay with you, just enter the dvips filename and let it run. If you want a PostScript file you can print locally, or if you want to use special ppages options, use the following:

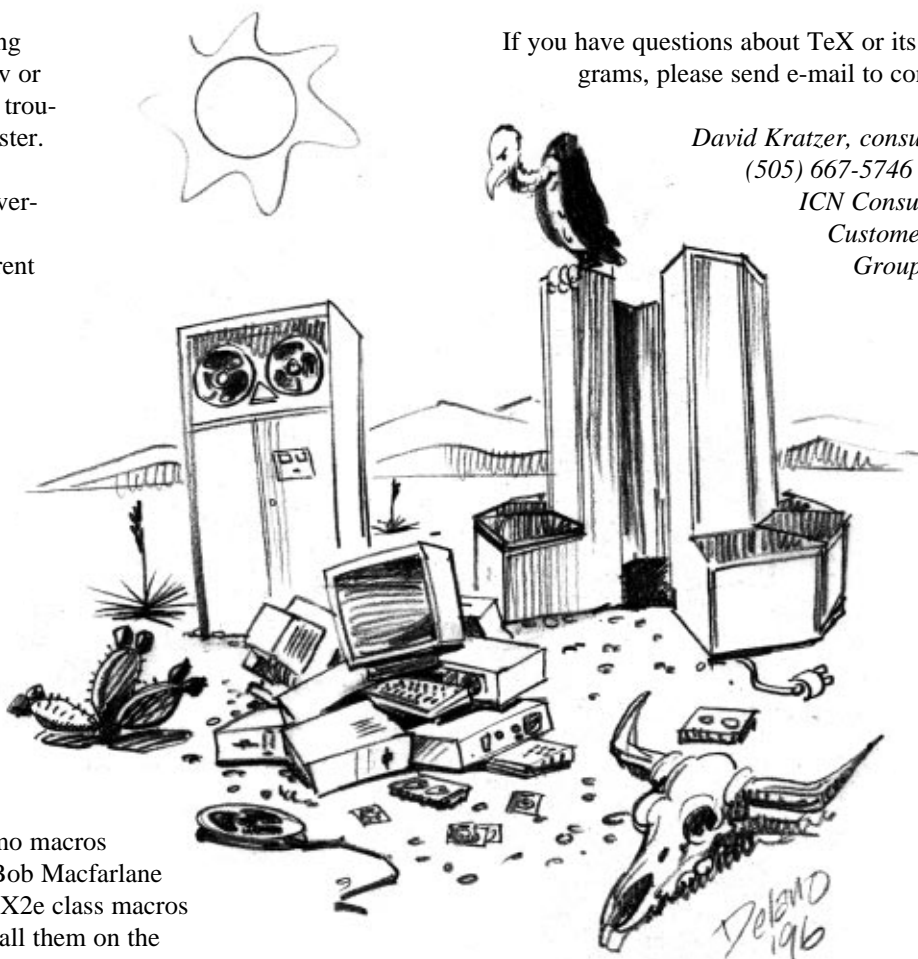
```
dvips -o filename.ps filename
```

This will process the file `filename.dvi` and leave `filename.ps` in your directory.

Man pages are available for most of the TeX related programs.

If you have questions about TeX or its related programs, please send e-mail to consult@lanl.gov.

David Kratzer, consult@lanl.gov,
(505) 667-5746
ICN Consulting Office
Customer Service
Group (CIC-6)



IBM XL High Performance Fortran Now Available on the Open Cluster

The IBM XL High Performance Fortran (HPF) for AIX (hereafter referred to as XL HPF) has been installed on the LANL Open Cluster and is ready for use. XL HPF is a compiler based on the HPF language specification. This specification provides a standardized set of extensions for Fortran 90 (F90) and Fortran directives that improve the performance of programs on parallel computer systems.

XL HPF provides programmers with a set of AIX-based tools, integrated with the AIX Common Desktop Environment (CDE), for the development of parallel applications running on AIX under RS/6000 SP systems and clusters of RISC System/6000 systems.

XL HPF is a native, highly optimized compiler that implements subset HPF features (with a few exceptions) and additional HPF features outside subset HPF. It supports F90 and F77 compilation and provides source code compatibility of existing F90 and F77 programs written with XLF for AIX.

XL HPF generates optimized executable objects with special focus on interprocessor communication and code optimization. It utilizes the following parallelization and optimization techniques:

- Extended optimization capabilities that allow analysis of data dependencies for better optimization and parallelization of DO-loops and F90 array language.
- Generation of in-line communication schedules for better performance.
- Message Vectorization—moves communication outside of loops to amortize the startup costs of sending messages.
- Collective Communication—identifies the high-level pattern of collective data movement, uses collective communication primitives like broadcast and reduction, and generates efficient send-receive code for special communication patterns.
- Elimination of Redundant Communication—detects and reduces redundant communications to increase performance.
- Coarse Grain Wavefronting—exploits pipeline parallelism.
- Mapping of Scalars—avoids unnecessary replication of scalar variables.
- Optimizes nearest-neighbor shift communication.
- Optimizes statement guards.

Enhancements available to F90 and F77 programs are as follows:

- Aggressive inlining of F90 array language to achieve performance comparable to F77 by eliminating unnecessary generation and copying of array temporaries.
- Loop structure transformed for improved data locality.

Common Desktop Environment (CDE)

XL HPF uses a new GUI based on the CDE in AIX Version 4.1.3 or later. CDE integration consists of an HPF application folder that is integrated within the CDE Application Manager. The HPF application folder contains icons representing the HPF tools and applications.

CDE integration of the HPF tools allows programmers to invoke the tools in a simple and consistent manner. The CDE desktop recognizes different types of files using a data type database. A data type identifies the files of a particular format and associates them with the appropriate applications. These associations mean that programmers don't have to remember command-line invocations of tools. In most cases, when a programmer double-clicks on a file, CDE will automatically launch the correct application to interpret that file's data.

The HPF application folder contains the following:

- Live Parsing Extensible (LPEX) editor,
- Program Builder,
- Debugger(xldb and pdbx),
- Command-line builder (xxlhp), and
- HPF on-line documentation.

LPEX Editor

The LPEX Editor is a language-sensitive, fully programmable editor that supports full F90 and some HPF functions. The LPEX Editor can be used to create and edit many types of text files, including program source and documentation. Using LPEX, developers can

- Use multiple windows to display several documents or to display more than one view of the same document,
- Dynamically configure LPEX to be a multiple-window or single-window tool,

- Select a block of text and move or copy it between documents,
- Cut and paste to a shell or other application, and
- Undo previous changes to a document.

Developers can customize and extend virtually every aspect of this programmable editor. LPEX is extended through dynamic link libraries. There is no proprietary extension language to learn. With the LPEX API, developers can write powerful extensions to the editor. In addition, LPEX provides a rich command language that developers can use to create or modify editor functions. Using LPEX, developers can

- Define their own fonts and colors,
- Modify the editor action key layout,
- Add menus to perform frequently used commands (menu definitions can be applied on a filename extension basis), and
- Write their own editor commands.

Program Builder

The Program Builder, a makefile generator that interprets XL HPF, manages the repetitive tasks of compiling, linking, and correcting errors in program source code. The Program Builder

- Provides a GUI to simplify the process of setting and saving compile and linker options,
- Supports error browsing from a list display (selecting a compile error in the list will position the programmer at the error in the source code within the LPEX Editor),
- Creates a “makefile” that is used by the AIX “make” command to construct and maintain programs and libraries, and
- Determines build dependencies by scanning the source code files for dependency information.

Debugger (xldb and pdbx)

The xldb is a GUI-based, serial full F90 symbolic debugger. The intuitive GUI allows programmers to control the execution of the program, examine and modify data (variables, storage, and registers), and perform many other useful functions.

The xldb debugger provides machine-level and source-level debugging. It is built around a set of core functions that let developers quickly and efficiently control execution and analyze data. With these core functions, developers can

- Display and change variables,
- Display and change storage,
- Display and change the processor registers,
- Display the call stack,
- Add and delete simple and complex breakpoints,
- Control the execution of multiple threads, and
- View source code as listing, disassembly, or mixed.

The pdbx is a parallel dbx debugger provided by Parallel Environment for AIX Version 2 and can be invoked via the HPF application folder under CDE. The pdbx debugger provides debugging support for Parallel Environment jobs.

Command-Line Builder (xxlhpfb)

The xxlhpfb is a GUI-based, command-line builder that interprets all options available for XL HPF. It simplifies the process of selecting compiler options and helps programmers to understand what each option does.

XL HPF Run-Time Environment (RTE)

XL HPF RTE provides a run-time environment to support processor communication, HPF library modules, I/O services, and utilities.

Open Enterprise

The XL HPF compiler was designed according to the HPF Language Specification Version 1.1 (released by the High Performance Fortran Forum on November 10, 1994) as understood and interpreted by IBM. XL HPF supports this specification with the following exceptions:

- ENTRY Statement;
- DISTRIBUTE Directive—CYCLIC(N), where N is not equal to 1;
- Multiple Processor Arrangements; and
- Internal I/O with subobjects.

In addition to the subset HPF features, XL HPF also provides the following additional HPF features for standardized HPF programming:

- PURE procedures;
- FORALL construct and statement;
- Storage and sequence association, including the SEQUENCE directive except for mapping of sequenced variables;
- HPF_LOCAL and HPF_SERIAL extrinsic kinds on subroutines and functions;
- HPF intrinsic procedures;
- Selected features of the HPF_LOCAL_LIBRARY module and HPF_LIBRARY module;
- Substantial support for F90 in HPF programs;
- Fortran 90 (ISO/IEC 1539-1991 and ANSI X3.198-1992);
- FORTRAN 77 (ANSI X3.9-1978 and Full ANSI F77);
- ISO 1539-1980, Programming Language, FORTRAN;
- FIPS PUB 69-1; and
- ANSI/IEEE Standard 754-1985 for binary floating point arithmetic.

Why Use XL HPF?

Parallel programming is a solution for programmers to achieve better throughput to solve complex problems. But parallel programs are difficult to write and maintain because there are many different types of parallel architectures and parallel programming environments. As a result, parallel programs tend to be non-portable across machines.

Data parallel programs can be written for RS/6000 SP systems or clusters of RISC System/6000 systems using XL Fortran for AIX with explicit message passing calls. This can be time consuming, error prone, and complicated in nature.

With the availability of the XL HPF compiler, programmers can write data-parallel applications using a standardized set of language extensions and directives in their program without needing to manage explicit message-passing calls. This helps to increase the productivity of programmers by allowing them

to focus on their application without spending time managing message-passing calls. In addition, HPF application programs are machine independent and can be ported to other parallel computer systems without major effort.

The XL HPF compiler allows programmers to write data-parallel programs without explicitly defining the number of processors. The HPF executable objects take advantage of the number of processors specified at execution time because recompilation of the HPF program is not required. If the number of processors is defined in the program, then the number of processors specified at execution time must match the number of processors defined in the program.

XL HPF exploits the RISC System/6000 processor architectures and is finely tuned to work with Parallel Environment for AIX by employing state-of-the-art optimization and parallelization technology.

*Bob Boland, wrb@lanl.gov, (505) 667-1729
Computing Group (CIC-7)*

*Doug Lora, dlora@lanl.gov, (505) 665-3321
IBM / Computing Group (CIC-7)*

Network Operations Center

Is your computer network on the blink?

Not receiving your e-mail?

**Can't reach the network through your
dial-up modem?**

**Want to know how to obtain an IP address
or host name for your computer or
network device?**

**Have questions about the Laboratory's
computer network backbone infrastructure
(routers, hubs, concentrators, etc.)?**

Let the Network Operations Center (NOC) in Network Engineering (CIC-5) answer your questions. Contact the NOC at 7-7423 between 7:00 a.m. and 5:00 p.m. Monday through Friday or by e-mail at noc@lanl.gov (after hours call 7-4585).

Research Library Training

The LANL Research Library provides training for using its specialized databases. Training sessions begin and end at times indicated below. Classes are free but you must pre-register by calling the Research Desk at 7-5809 or sending e-mail to library@lanl.gov. Special classes and orientations can also be arranged.

Date	Time	Subject Matter
11/5/96	1:00 -1:30 p.m.	Commercial Information for Patent Applications
11/6/96	1:00 -1:30 p.m.	Finding Addresses and Phone Numbers on the WWW
11/7/96	1:00 -1:30 p.m.	SciSearch Alerting Service
11/7/96	2:00 - 4:00 p.m.	Information Sources on the Internet via WWW
11/12/96	1:00 -1:30 p.m.	Federal Regulations on the Internet
11/14/96	1:00 -1:30 p.m.	SciSearch at LANL—At your desktop!
11/19/96	1:00 -1:30 p.m.	Finding Secret Information (Q-Clearance Required) (Held in the Report Collection)
11/20/96	11:00-11:30 a.m.	MELVYL (U of CA specialized databases)
11/20/96	1:00-1:30 p.m.	Finding Addresses and Phone Numbers on the WWW
11/21/96	1:00-1:30 p.m.	Search Engines for the WWW
11/21/96	2:00 - 4:00 p.m.	Information Sources on the Internet via WWW
11/26/96	1:00-1:30 p.m.	Engineering Index—At your desktop!

Labwide Systems Training

The Customer Service Group (CIC-6) offers training for users of Laboratory information systems. The CIC-6 courses offer training for a variety of personnel including property administrators, group secretaries, training coordinators, budget analysts, group leaders, or anyone needing to access training records, property records, costs, employee information, travel, chemical inventories, etc. Refer to the table below and on the following pages for specific information about courses currently offered.

Course Registration

You must have a valid ICN password before taking any of the courses shown in the table. To register for a course, call the CIC-6 Training, Development, and Coordination section at 667-9559 or access our Web page. From the LANL home page, look under "Services/Computing at LANL/Training" or enter the URL:

<http://www.lanl.gov:8010/computer-information/cic6/teampage.html>

Course Title	Date	Time	Cost	Course Number
Employee Development System - Basic Training (EDS I):	11/6/96	8:30 – 12:00	\$350	Course #5289
The course provides hands-on instruction to request course enrollment, use the on-line course catalog, retrieve training transcripts, and assign EDS authorities. The student will learn to create courses, add students to the courses, and generate several training reports.				
Employee Development System - Training Plans (EDS II):	11/20/96	8:30 – 12:00	\$350	Course #7155
Participants receive hands-on instruction to create and maintain training plans, assign assignment codes, and generate training plan reports. Attendees must have prior training in the Employee Development System (course #5289).				
Eudora Electronic Mail	11/13/96	1:30 – 3:30	\$175	Course #9762
This class is a hands-on class that teaches the participant how to use Eudora software to create, send, receive, and edit electronic mail messages. In addition to these procedures, the participant will learn what related settings mean and how to configure the system to meet his or her individual needs.				
Data Warehouse Basics	11/15/96	8:30 – 10:30	\$175	Course #11961
Students will receive hands-on training to generate standard reports and make quick queries from information in the data warehouse, a real-time collection of data tables from Laboratory financial, time-reporting, and personnel systems.				
Data Warehouse/ Financial Reporting	11/15/96	8:30 – 12:00	\$350	Course #11960
Students will receive hands-on training to generate standard financial reports and make on-line queries from information in the "data warehouse," a collection of data from Laboratory budgeting, accounting, and time-keeping systems.				
HTML Basics	11/19/96	8:30 – 12:00	\$350	Course #11605
Students will gain a basic understanding of HTML (Hypertext Markup Language), the language for the World Wide Web. Topics covered will be commands and standards, creating and editing documents, and authoring programs.				

Course Title	Date	Time	Cost	Course Number
HTML Tables	December	8:30 – 12:00	\$350	Course #11959
Students gain basic understanding of how to create various tables in HTML and new tags in HTML 3.0. Netscape-specific tags are also identified for clarity. Prerequisite: HTML Basics (Course #11605) or permission of the instructor.				
Introduction to the Internet: Beginning Netscape	December	1:30 – 3:30	\$175	Course #10961
Students gain basic understanding of the Internet and the World Wide Web and the use of Netscape as a browser to surf the Net. Topics covered are both Laboratory sites and open sites, along with practical uses of the Internet.				
Lotus Notes 4.0	December	1:30 – 5:00	\$350	Course #9917
This class provides hands-on instruction for Mac and PC users to use Lotus Notes software to create and send E-mail memos; fax documents; search databases; create filters, nicknames, banners, and doclinks; set defaults; and use multiple address books. In addition, participants learn how to use the memo, meetings, and discussion databases.				
On-Line Forms	December	3:30 – 5:00	\$175	Course #9756
Participants will learn to use Netscape software to access Lab-wide information and forms. Using Jetform Filler software, participants will access, complete, and print forms such as the "ICN Validation Request," "Visitor Request for Unclassified Visits to Security Areas," and "Request for Quotation."				
Purchase Card System	11/4/96	1:30 – 2:30	\$175	Course #11924
Students will learn to reconcile monthly statement of account, submit reconciled statement of account for approval, print statement of account for audit records, and delegate reconciliation authority. Prerequisite: PCS Overview. Call Ruby O' Rear at 665-4523.				
Reporting with Infomaker	December	8:30 – 5:00	\$650	Course #11054
Hands-on training to query data and develop ad hoc, or non-standard, reports from the LANL data warehouse using Infomaker software.				
Time and Effort System (GUI)	11/14/96	8:30 – 10:00	\$175	Course #11018
The student will learn how to enter attendance, amend attendance, approve attendance, and submit exception and approval reports. Time codes and associated policies will also be discussed. In addition, the student will learn how to use the Information Manager utility to view and print reports.				
Travel	11/21/96	8:30 – 11:30	\$350	Course #12091
Hands-on training to submit and approve travel requests and expenses in the new Travel System which replaces the TRIPS on-line system and the post-travel expense worksheets.				

Vendor Computer Training

The Customer Service Group (CIC-6) supports vendor training in technical computing areas such as programming languages, system administration, networking, and World Wide Web development tools. The support provided by CIC-6 can be as limited as providing the appropriate facilities for a specific group or as extensive as coordinating training functions such as system administration, vendor acquisition, EDS administration, and class facilitation. The table below lists classes that are either currently being offered or are available on request. An expanded list of classes that are potentially available can be viewed on the Internet at

<http://www.lanl.gov:8010/computer-information/ComputerTraining/Vendor.html>

To request registration in any vendor course or for general assistance with vendor training, please contact the CIC-Division Vendor Training Coordinator at (505) 667-9399 or send e-mail to cic6-train@lanl.gov.

*Cost per student will vary depending on the total number of students enrolled in the class.

Course Title	Date	Time	Cost	Course Number
C Programming (Beginning)	Available on Request (5 days)		\$1200–\$1700*	3996
Prerequisite(s): An understanding of and useful skills in a high-level programming language. A current ICN password is required. Topics Include: Introduction and Fundamentals; Basic Semantic Constructs - Getting; Base Level I/O With C; The Preprocess-Compilation Environment; Operators, Data Types, and Storage Classes; Control Flow Constructs; Conditional Constructs; Higher-Level Data Constructs in C; File I/O; UNIX Software Tools and POSIX System Calls.				
C Programming (Advanced)	11/18–22/96	8:30–5:00	\$1200–\$1700*	4777
Prerequisite(s): Useful skills and experience with the C Programming. A current ICN password is required. Topics Include: Data Structures, Algorithms, and OOP; An Advanced Clinic for C ; The ANSI C Recommendation X3.159; C and ANSI C War Stories; The Data Structure and the Assessment of Algorithms; Arrays; Structures; Unions; Stacks; Queues; Linked Lists; Recursive Functions; Binary Trees; Hashing; File Organizations Using the C Runtime Library; Standard Interprocess Communication Mechanisms; and An Introduction and Overview of AT&T's C++ 3.0.				
C++ for Experienced Programmers	12/16–20/96	8:30–5:00	\$1200–\$1700*	9050
	2/24–28/97	8:30–5:00	\$1200–\$1700*	9050
Prerequisite(s): Excellent C Language programming skills. Topics Include: Major Differences and Additions to ANSI C; Building C++ Classes; Introduction to Text I/O with C++; Function Overloading; Single Inheritance; Virtual Functions; Multiple Inheritance; Operator Overloading; Creating, Initializing and Assigning Objects; Passing and Returning Objects; Templates, Parameterized Functions and Classes; C++Stream I/O with the File System; and C++ Course Summary.				
Managing Internet Mail: Setting Up and Troubleshooting Sendmail and DNS	Available on Request (3 days)		\$1300-\$1800*	
Prerequisite(s): General knowledge of Unix system and network administration as well as experience with sending and receiving Internet electronic mail. Topics Include: Introduction to Using Electronic Mail; Theory of sendmail Operation; Understanding the sendmail.cf File; Address Rewriting Rules; Debugging sendmail; Understanding the Function of Sub-Domains in a Complex Mail Network; Setting Up Mail Sub-Domains and Mail Routing Hubs; Mail eXchanger (MX) Records and Mail Delivery in the Internet; Setting Up the Domain Naming System; Sendmail 8 - The Next Generation; Automatic Creation of sendmail.cf Files for Sendmail 8; and Verifying and Debugging sendmail.cf Files Generated by the sendmail Compiler.				

Course Title	Date	Time	Cost	Course Number
Object-Oriented Analysis and Design	1/21–24/97	8:30–5:00	\$1200–\$1700	
	1/27–30/97	8:30–5:00	\$1200–\$1700	
Prerequisite(s): Familiarity with fundamental programming concepts (data structures, types, control flow selection, iteration, etc.). Prior experience in systems or software analysis and/or development is useful but not required. Topics Include: Introduction to Object-Oriented Technology; The Object Model; OOAD Comparisons; Object-Oriented Analysis and Design I and II; Object-Oriented Analysis and Design Workshop; Object-Oriented Analysis/Design Methodologies; Object-Oriented Tools; Case Study: Texas Instruments; and Management Issues.				
Perl Programming	Available on Request (1–3 days)		\$500–\$700/day*	8095/8093
Topics Include: Describes the programming language that occupies the niche between shell and C Programming; syntax and semantics; data types; operators, control flow, regular expressions, and I/O facilities; the Perl debugger.				
Perl Programming for the WWW	Available on Request (2–3 days)		\$500–\$700/day*	
Prerequisite(s): Programming skills with a light background in Perl and HTML. Topics Include: On-line Resources; Server Configuration; Permissions; Setuid Issues; Tainting; Safe Perl; Data Security; OO Programming; Web Modules; CGI Programs; CGI.pm; What Went Wrong?; CGI Template; Using Forms; Form Template; Input Widgets; Submit Widgets; Reset Widgets; Sample Form; Password Fields; Text areas; Hidden Fields; Checkboxes; Radio Boxes; Popup Menus; Lisboxes; Image Maps; Random Links; libwww Modules; Sending Mail; Shopping Carts; Database Access; and Advanced Topics.				
SGI System Administration (Beginning)	Available on Request (5 days)		\$1800–\$2300*	7993
Prerequisite(s): Familiarity with using Silicon Graphics IRIS workstations and system administration procedures on other open system platforms. Topics Include: The Role of the System Administrator; Set Up and Configuration of an IRIS Workstation or Server; Supporting a Group of Silicon Graphics Users; System Security Maintenance; Backups and Recoveries; Configuration of Disk Drives; System Installation and Application Software; Attaching Terminals and Printers; Modifying the system Start Up and Shut Down Sequences; Automating Administrative Procedures; and Performing Basic System Troubleshooting.				
SGI Network Administration	Available on Request (5 days)		\$1800–\$2300*	11690
Prerequisite(s): Completion of Silicon Graphics System Administration (Beginning) course or equivalent knowledge and experience. Topics Include: Networking Fundamentals; Network Configuration; Network Troubleshooting; Resource Management with Network; Information Services; Domain Management with Domain Name System; Electronic Mail with Sendmail; Remote File Sharing with Network File System & Automounter; Network Performance Monitoring; and Network Security.				
SGI System Administration (Advanced)	Available on Request (5 days)		\$1800–\$2300*	11689
Prerequisite(s): Completion of Silicon Graphics System Administration (Beginning) course or equivalent knowledge and experience. Topics Include: System Error Monitoring; Kernel Reconfiguration and Debugging; System Monitoring Tools; Process Management; MultiProcessor CPU Management; Memory Management and Tuning; Swap Management and Tuning; Disk Management and Tuning; XPS Filesystem Management; and System Security Concepts.				

Course Title	Date	Time	Cost	Course Number
Solaris 2.X System Administration (Beginning)	Available on Request (5 days)		\$1600–\$2000*	7477
Prerequisite: Knowledge of Unix commands and an editor. Topics include: Custom installation of Solaris2.X server; Add peripheral devices; Use format utility to display partition information; Compress and send binary files; Change system run levels; Add startup files for additional services; Add and remove software packages; Configure terminals and modems; Administer disks and file systems; Discuss basic networking concepts; Configure NFS to support the client-server environment; Use the automounter; Add and remove diskless clients; Back up and restore file systems; Perform basic recovery and troubleshooting procedures; Configure and administer the NIS+ environment.				
UNIX (Beginning)	11/18–22/96	8:15–12:00	\$738	5267
Prerequisite(s): Familiarity with a UNIX workstation. Topics Include: Overview of the Workstation Environment; Getting Started; The UNIX File System; Manipulating Files; Customizing Your Environment; The C-Shell; Editing and Writing with vi; Using the Network; Discussing NFS and NIS; Using Basic System Status Commands; Startup and Shutdown Procedures; Using tar.				
Windows NT Workstation and Server	Available on Request (5 days)		\$1600–\$1900*	
Prerequisite(s): This course is valuable for personnel who are evaluating or migrating to Windows NT. It benefits system and network administrators, other support personnel, programmers, and users from Windows, Unix, OS/2, or VMS backgrounds. Topics Include: Introduction to Windows NT; System Overview and Security; Network Configuration Options; Installation; Server Choices; User Administration and Security; Files and Printers; Built-in Network Support; Configuration Options; Using Setup; Data and Disk Management; The Registry; Troubleshooting; and Optimization and Performance.				

INTEGRATED COMPUTING NETWORK (ICN) VALIDATION REQUEST

To access ICN Computing resources, please complete all parts of this form that apply to you, including "Special Requirements."

Mail your completed application to:
ICN Password Office (PWO)
Mail Stop: B271
Los Alamos National Laboratory
Los Alamos, NM 87545

If you have **questions**: Call: (505) 665-1805
E-mail: validate@lanl.gov

All Laboratory computers, computing systems, and their associated communication systems are for official business only. By completing this request, users agree not to misuse the ICN. The Laboratory has the responsibility and authority to periodically audit user files.

Owner Information

Z-Number (If you have one)	PWO Use Only	Name (last, first, middle initial)
LANL Group	LANL Mail Stop	Citizenship (Foreign National see "Special Requirements-Foreign National")
Phone Number	Cost Center	Program Code

<p>Check LANL affiliation:</p> <p><input type="checkbox"/> LANL employee</p> <p><input type="checkbox"/> Contractor _____ (specify contract company)</p> <p><input type="checkbox"/> Consultant, VSM, associate</p> <p><input type="checkbox"/> External user _____ (specify employer)</p> <p><input type="checkbox"/> Other (specify) _____</p>	<p>Send password / smartcard to:</p> <p><input type="checkbox"/> Mail Stop or <input type="checkbox"/> Mail to address indicated below</p> <p>Name / Organization</p> <hr/> <p>Address</p> <hr/> <p>City, State, Zip Code</p> <hr/>
---	--

Access Check access method and needed partitions:


Access method:	<input type="checkbox"/> ICN Password	<input type="checkbox"/> Smartcard	<input type="checkbox"/> Both
<input type="checkbox"/> Open partition (e.g., email systems, open machines)			
<input type="checkbox"/> Administrative partition (e.g., IA [BUCS, Stores, Travel], IB [EIS, FMIS, PAIRS]) If you are not a Q-cleared LANL employee, see required steps in section "Special Requirements-Administrative Partition," unless you already have Administrative access with an ICN password.			
<input type="checkbox"/> Secure partition (i.e., secure machines) Indicate level(s) of data to be processed:			
<div style="display: flex; justify-content: space-between;"> <div style="width: 40%;"> <input type="checkbox"/> Unclassified <input type="checkbox"/> Secret </div> <div style="width: 55%; border: 1px solid black; padding: 5px;"> I certify this person does require secure access: <div style="display: flex; justify-content: space-between; margin-top: 10px;"> <div>_____ Manager Signature (Group Leader or above)</div> <div>_____ Date</div> </div> </div> </div>			

NOTE: A Q-clearance is required. All classified computing must be performed within the Secure environment.

PWO Use Only

<input type="checkbox"/> New <input type="checkbox"/> Change	Clearance Status	Processed	Lv	Smartcard Serial #
Comments:				

Form 1848 (1/95) Supersedes previous versions (rev. 10/20).

Continue 

Special Requirements

Administrative Partition

(U.S. Citizens Only)

Lab-Wide Systems (e.g., IA [BUCS, Stores, Travel], IB [EIS, FMIS, PAIRS])

☐ Under 18 years of age

If you need to access Administrative systems, your group leader must provide a memo accepting responsibility for your actions and justifying your need for access. This memo is to accompany all forms taken to the security briefing (see "Contractor or Non-Q-Cleared" section below). You may not access the Secure Partition.

☐ Contractor or Non-Cleared

Phone (505) 667-9444 to obtain Access Authorization packet.

Phone (505) 667-9153 to schedule a security briefing.

Bring all forms including this ICN Validation Request to the security briefing for approval.

Security Briefing Approval Signature

Date

☐ Foreign National

Attach a copy of Form 982 (REQUEST FOR UNCLASSIFIED VISIT OR ASSIGNMENT BY A FOREIGN NATIONAL) with all approval signatures. Be sure Box #11 of Form 982 is completed. If you are not a visitor/assignee under a LANL/DOE approved Visit / Assignment Request, attach written justification from your host Division Director describing your need to access the ICN.

Authorization (required)

Print Manager Name (Group Leader or above)	Manager Z-Number	Group
Manager Signature (Group Leader or above)	Mail Stop	Date

If you are NOT a LANL employee, obtain your LANL contact's signature in addition to the contact's manager's signature.

NOTE: LANL contacts are regular Laboratory employees. Contacts are responsible for obtaining annual re-authorizations, forwarding renewals, and notifying the ICN Password Office of changes in user or contact status.

Print LANL Contact Name	Contact Z-Number	Phone Number	Group
LANL Contact Signature	Mail Stop	Date	

Reader Feedback

Feedback helps us to provide a document that responds to the changing needs of its readership. If you have comments or questions about this publication, please let us hear from you. We have reserved the back of this form for that purpose. We also accept articles for publication that are of interest to our readers. Contact the managing editor for more information. This form is also used for new subscriptions, deletions, or changes. Instructions are on the back. If you prefer to contact us by E-mail, send your comments and/or subscription request to finney@lanl.gov.

Do Not Staple
Fold on This Line First



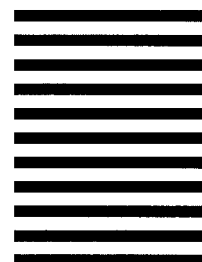
NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 88 LOS ALAMOS NM

POSTAGE WILL BE PAID BY THE ADDRESSEE

MAIL STOP B251
ATTN: MIKE FINNEY, MANAGING EDITOR
CUSTOMER SERVICE GROUP (CIC-6)
LOS ALAMOS NATIONAL LABORATORY
PO BOX 1663
LOS ALAMOS NM 87544-9916



Do Not Staple, Seal with Tape
Fold Here

cut along dashed line

This image shows a single sheet of white paper with horizontal ruling lines. The lines are evenly spaced and run across the width of the page. There are no margins, text, or other markings on the paper.

BITS is published by Los Alamos National Laboratory. If you would like to be added to or deleted from our mailing list, please check the appropriate line, complete the form below, and mail us the form.

_____ Delete my name from the **BITS** mailing list.

_____ Change my name/address as indicated below.

Date

Mail Stop

Organization

Zip

Employee Z#

INDEX

Keywords	Title of BITS Article	Date	Page
<i>BITS</i>	<i>Announcing a Special Edition of BITS: Introduction to Computing at Los Alamos</i>	May '96	1
	<i>Special Edition of BITS Available On-Line</i>	July '96	10
<i>CFS (Common File System)</i>	<i>Scripts for Copying Filetrees between CFS and UNIX: Wood Man, Spare That Tree!</i>	Mar. '96	10
	<i>Remote-Backup Service Provided by CFS</i>	Dec. '95	2
<i>CGI (Common Gateway Interface)</i>	<i>CGI Security</i>	Mar. '96	6
<i>Challenge</i>	<i>Challenge to Start Seventh Year</i>	Sept. '96	4
<i>CIC (Computing, Information, & Communications)</i>	<i>Holiday Schedule for CIC Computing Services</i>	Dec. '95	2
	<i>CIC Helps U.S. Postal Service Provide Secure Electronic Commerce for American Citizens</i>	Feb. '96	1
	<i>FY96 Rates for CIC Products and Services</i>	Feb. '96	3
	<i>New CIC Annual Report Available On-Line</i>	July '96	10
	<i>CIC Implements New Recharge Processing System</i>	Sept. '96	1
<i>CIC-4</i>	<i>CIC-4 Awards External Paging Contract and Expands Local Paging Capabilities</i>	Aug. '96	5
<i>CIC-6</i>	<i>The CIC-6 Vendor Training Program</i>	Apr. '96	5
<i>CIC-12</i>	<i>CIC-12 Employee Responds to "Unthinkable"</i>	Mar. '96	1
<i>CIC-17</i>	<i>CIC-17 Offers Highlight Color Printing</i>	Apr. '96	12
<i>Cluster</i>	<i>New AIX Parallel Environment and HPF Products Now Available on the Open IBM Cluster</i>	Dec. '95	11
	<i>Introducing the Cluster Web Pages</i>	Feb. '96	18
	<i>Configuration Changes on the Open Cluster: Announcing the New SP2</i>	Mar. '96	11
	<i>Software Currently on the Open and Secure Clusters</i>	Mar. '96	14
	<i>Cluster Tutorials on the Web</i>	Aug. '96	8
	<i>Maple on the Cluster</i>	Aug. '96	11
<i>Cray</i>	<i>Cray Programming Environment 2.0</i>	Feb. '96	2
	<i>Cray C++ Programming Environment 2.0</i>	Mar. '96	3
	<i>Cray CF90 Programming Environment 2.0 for PVP and SPARC Systems</i>	Apr. '96	6
	<i>Cray Programming Environment 2.0 Tools</i>	May '96	2
	<i>The New Opportunity Scheduler for Crays</i>	May '96	8
<i>Database</i>	<i>ABI/INFORM Business Periodicals Database Now Available at Your Desktop</i>	Dec. '95	3
	<i>NTIS Research Database at Your Desktop</i>	July '96	4
<i>Desktop</i>	<i>What's on Your Mac's Desktop?</i>	Dec. '95	8
	<i>Laboratory Desktop Software Standards</i>	Apr. '96	1
<i>E-mail Attachments</i>	<i>How to Get Unattached to E-mail Attachments Part 1: A Look at the Macintosh</i>	May '96	10
	<i>How to Get Unattached to E-mail Attachments Part 2: A Look at the PC</i>	June '96	8
	<i>Don't Get too Attached to Your [E-mail] Attachments</i>	Oct. '96	8
<i>Enterprise Information Systems</i>	<i>Enterprise Information Systems</i>	Oct. '96	6
<i>Environmental Cleanup</i>	<i>New Computer Programs Support Environmental Cleanup</i>	Dec. '95	1
<i>Eudora</i>	<i>How to Point the Eudora Finger in the Right Direction</i>	Feb. '96	14
<i>HPD (Heterogeneous Parallel Debugger)</i>	<i>HPD: Heterogeneous Parallel Debugger</i>	Nov. '95	2
<i>HTML (HyperText Markup Language)</i>	<i>Standard HTML Reaches More People and Saves Time</i>	Dec. '95	7
	<i>Tips on Writing HTML <TABLE>s</i>	Feb. '96	10
<i>ICN (Integrated Computing Network)</i>	<i>Massively Parallel Supercomputing in the Secure ICN</i>	Feb. '96	1
	<i>ICN Host Status Now Available on the Web</i>	Apr. '96	17
	<i>Questions and Answers from the ICN Consulting Office</i>	June '96	2
	<i>ICN Password Office Provides FAQ Web Page</i>	Oct. '96	7
<i>Information Architecture (IA)</i>	<i>IA Project Proposes Web Publication Guidelines</i>	Dec. '95	4
<i>Internet</i>	<i>Responsible Use of the Internet</i>	July '96	1
<i>ISDN (Integrated Services Digital Network)</i>	<i>Orders for ISDN Suspended</i>	July '96	4
<i>Lab-Wide Systems</i>	<i>Questions and Answers for Lab-Wide Systems Users</i>	Apr. '96	14

Keywords	Title of BITS Article	Date	Page
<i>LSF (Load Share Facility)</i>	<i>Life in the Fast Lane: Using LSF to Cut Down on Cluster Computing Time</i>	<i>June '96</i>	<i>3</i>
<i>Macintosh</i>	<i>Four Macintosh Shareware Programs That Will Make Scientists Happy</i>	<i>Sept. '96</i>	<i>9</i>
<i>Medicare</i>	<i>Computer Sleuths Hunt for Medicare Bandits</i>	<i>Sept. '96</i>	<i>2</i>
<i>Mercury</i>	<i>Mercury Open-Secure ICN File Transfer Service Now Available</i>	<i>Oct. '96</i>	<i>4</i>
<i>MPI (Message Passing Interface)</i>	<i>Getting the Most out of MPI</i>	<i>July '96</i>	<i>5</i>
<i>MIME (Multipurpose Internet Mail Extensions)</i>	<i>A Closer Look at MIME</i>	<i>July '96</i>	<i>11</i>
<i>Network</i>	<i>Statistics on Network and Telephone Services at LANL</i>	<i>May '96</i>	<i>4</i>
<i>OAG (On-line Airline Guide)</i>	<i>OAG Service Announcement</i>	<i>Sept. '96</i>	<i>5</i>
<i>PAGES (Print And Graphics Express Station)</i>	<i>PAGES Hardware Status Now Available on the Web</i>	<i>Apr. '96</i>	<i>13</i>
<i>Password</i>	<i>On-Line Semiannual Passwords for Secure Users</i>	<i>Apr. '96</i>	<i>5</i>
<i>Picking an On-Line Name</i>	<i>Picking an On-Line Name for Yourself, Team, or Project</i>	<i>Mar. '96</i>	<i>5</i>
<i>PVM (Parallel Virtual Machine)</i>	<i>PVM: Easier to Use</i>	<i>Apr. '96</i>	<i>8</i>
<i>Research Library</i>	<i>Research Library's Monthly Electronic Newsletter</i>	<i>Aug. '96</i>	<i>8</i>
<i>Secure Unclassified</i>	<i>Good-Bye Secure Unclassified</i>	<i>Oct. '96</i>	<i>5</i>
<i>Society and the Future of Computing '96</i>	<i>Society and the Future of Computing '96: Call for Participation</i>	<i>Mar. '96</i>	<i>4</i>
<i>Telephone Services</i>	<i>FTS 2000 Provides Savings for LANL Telephone Services</i>	<i>Aug. '96</i>	<i>6</i>
<i>World Wide Web (WWW or Web)</i>	<i>The LANL Web Mistress</i>	<i>Dec. '95</i>	<i>6</i>
	<i>Web Browsers and Helper Applications</i>	<i>Feb. '96</i>	<i>16</i>
	<i>CIC Pilots Web Design Guidelines and Offers Web Classes</i>	<i>Mar. '96</i>	<i>9</i>
	<i>Creating Effective Web Page Backgrounds</i>	<i>Apr. '96</i>	<i>15</i>
	<i>Mathematical Reviews Available On-line through the Research Library Web Page</i>	<i>May '96</i>	<i>3</i>
	<i>Copyright and the World Wide Web</i>	<i>May '96</i>	<i>5</i>
	<i>CIC-1 Web Team Designs and Revises Web Pages</i>	<i>June '96</i>	<i>1</i>
	<i>Barnstorming the Web</i>	<i>June '96</i>	<i>5</i>
	<i>Images on the Web: Some Tips</i>	<i>Aug. '96</i>	<i>1</i>
	<i>Finding Phone Numbers and Addresses via the Web</i>	<i>Sept. '96</i>	<i>5</i>
	<i>Images on the Web: More Tips</i>	<i>Sept. '96</i>	<i>6</i>
	<i>Mariachis Weave Beethoven? More Tips and Tricks [on the Web]</i>	<i>Oct. '96</i>	<i>1</i>
<i>Windows 95</i>	<i>Windows 95 SLIP/PPP Installation</i>	<i>Apr. '96</i>	<i>18</i>

Produced by the Computing, Information, and Communications (CIC) Division

Managing Editor: Mike Finney (667-2241 or finney@lanl.gov)

Design: Gloria Sharp and Mike Finney

Illustration: Dave Delano

Printing: Media Group (CIC-17)

Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by the University of California for the United States Department of Energy under contract W-7405-ENG-36.

All company names, logos, and products mentioned herein are trademarks of their respective companies. Reference to any specific company or product is not to be construed as an endorsement of said company or product by The Regents of the University of California, the United States Government, the U.S. Department of Energy, nor any of their employees. The Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; therefore, the Laboratory as an institution does not endorse the viewpoint of a publication or guarantee its technical correctness.

Los Alamos
NATIONAL LABORATORY

Los Alamos

NATIONAL LABORATORY

Los Alamos, New Mexico 87545

BITS is published monthly to highlight recent computing and communications activities within the Laboratory.

We welcome your suggestions and contributions.

BITS can be accessed electronically via Web browsers such as Mosaic and Netscape. Enter the following URL:

<http://www.lanl.gov/Internal/divisions/cic/publications.html>

LALP-96-11 (10-96)

Nonprofit
organization
US Postage

PAID

Los Alamos, NM
Permit No. 107